Different Methods for Image Inpainting

Mustermann Max
Computer Vision Lab
TU Wien
Vienna, Austria
e000000@student.tuwien.ac.at

I. Introduction

The goal of image inpainting is to restore parts of an image, in such a manner, that a viewer can not detect the restored parts. One application of image inpainting is to retouch damaged parts of a digital picture. Before the inpainting process is started, the user defines a binary mask for the image, which marks the region that should be restored. The following image inpainting task is automated and needs no further user interaction. The term digital image inpainting was coined by Bertalmio et al. [1]. The authors of this paper suggest a method for inpainting that is based on Partial Differential Equations (PDE). The underlying idea of the suggested method is to smoothly complete isophote lines, arriving at the border of the region that should be inpainted, from the out- side of the border to the inner region. Two drawbacks of image inpainting methods based on PDE are that they only perform well on small inpainting regions and that they are not able to fill in texture [2].

Criminisi et al. present a method - called exemplar-based image inpainting [2] - that uses the main idea PDE and is able to fill in texture. This heuristic approach also makes it possible to fill bigger regions.

Instead of using a heuristic approach, Roth et al. propose an algorithm - called Fields of Experts (FoE) [3] - that is based on probability theory. The authors use a model - which is trained on an image database - to describe the continuity of image features (like edges for example). With this model it is possible to carry out the image inpainting task.

In the following sections the work by Criminisi et al. and Roth et al. will be examined in more detail. Later on results of both algorithms will be given.

II. Exemplar-based Image Inpainting

With the use of texture synthesis researchers it is possible to overcome the drawbacks of PDE based inpainting methods. For image inpainting the most popular texture synthesis technique is called exemplar-based synthesis. [4] This method fills the inpainting region with a texture, generated by texture patches from the surrounding areas. Criminisi et al. suggest a method - called exemplar-based image inpainting [2] - that combines the strengths of PDE based methods and exemplar-based techniques.

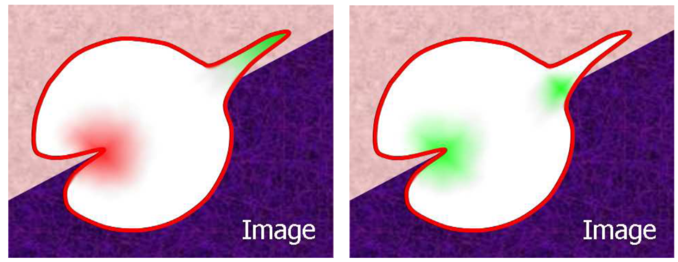

(a) Visualization of the confidence term. (b) Visualization of the data term.

Fig. 1: Assignment of the fill priority. (a) The graphic shows how the confidence is assigned: Pixels belonging to the green region possess more confidence than pixels belonging to the red area, since the pixels in the green area are surrounded by more known pixels. (b) The data term gives pixels that belong to edges, a higher priority. (Those pixels are painted green in the figure.) [2]

A. Fill Order

The main focus of PDE-based image inpainting methods lies on the preservation of linear structures (e.g. edges). The authors take this circumstance into account, by defining a fill order, which gives pixels belonging to edges a higher priority, than pixels belonging to homogenous regions. The authors of the paper call this kind of priority the data term.

The priority of one pixel is also determined by its confidence, which is a measurement of the reliability of the information surrounding the pixel: The confidence is directly proportional to the number of the pixels in the neighborhood that have been known from the beginning, or that have already been filled. This measurement is called the confidence term. The overall priority of a pixel is the product of its data term and its con- fidence term. In Figure 1 an example for the priority assignment is shown.

B. Filling Process

The subsequent filling process starts with the pixel with the highest priority. The algorithm searches an image patch in the known (or already filled) image region that has the highest similarity to the patch surrounding the observed pixel. The similarity is measured with the sum of the squared difference of the already filled pixels in the two patches. Once the patch with the highest similarity has been found, the color value of the pixel at the center of the patch is assigned to the observed pixel.

III. Field of Experts

While the previously described approach for image inpainting is a greedy algorithm, the second approach - that is described in the following - is based on probability theory. Since a detailed description of the underlying theory is far beyond the scope of this abstract, this section contains only the basic concepts of the Fields-of-Experts (FoE) algorithm. For a detailed description of the algorithm the reader is referred to the publication of Roth et al. [3].

A. Calculation of the prior

The underlying theory of the FoE algorithm is Bayesian inference, which makes it possible to calculate the prior of the image. The prior is the undamaged version of the input image and it is estimated with the observed (damaged) image. The prior is calculated with the usage of Markov Random Fields (MRF). The basic idea behind MRF is that an image is a random field. The value of one pixel of the image only depends on the value of its neighboring pixels. The probability that the pixel has a certain value can be modelled with a special distribution that depends on a chosen function - called the energy function. The prior is afterwards found by maximising the probability of the pixel values or minimizing the energy function.

It is very essential for the success of MRF algorithms, to find an appropriate energy function. In the past the parameters for this function were typically hand-crafted. Roth et al. instead propose an algorithm that is able to learn those parameters by training a model on an image database. The trained model consists of several convolution filters that serve as the desired parameters. The prior can be computed, as the product of the filter responses on the observed image.

B. Image Inpaintings with Fields-of-Experts

For the inpainting task only the pixels in the region that should be inpainted are modified. In an iterative process a gradient ascent procedure is performed, to find the image prior (which means filling the unknown region). The algorithm starts with the computation of the prior of the observed image. The gradient of the prior is afterwards added to the inpainting region of the observed image. This new image is used for the next iteration, where its gradient is computed and added to the inpainting region. The algorithm stops the inpainting task after a predefined number of iterations.

IV. Results

(a) (b)

Fig. 2: (a) Original Photograph. (b) The resulting inpainted image. The inpainting region was defined manually. [2]

(a) (b)

Fig. 3: (a) Original image; the red text is the inpainting region. (b) The resulting inpainted image. [3]

In Figure 2 two resulting images of the exemplar- based image inpainting method are given. The images show that the algorithm is able to remove large objects from images with a heterogenous background.

Figure 3 shows two images that were in- painted with the FoE algorithm. The authors compare their method with the PDE-based image inpainting by Bertalmio et al. [1] and find out, that the continuity of edges is better preserved with their algorithm. Unfortunately Criminisi et al. used different image for their tests, than Roth et al. did, and therefore no direct comparison between the two described inpainting can be made.

V. Conclusion

In this abstract two different approaches for image inpainting were presented. The first one is a greedy algorithm that copies texture parts from the known image region into the unknown region. The algorithm pays special attention to the preservation of edges. The second algorithm is based on statistical theories and trains a model on an image database. With the model it is possible to calculate the prior of an observed image. One limitation of the second approach is that it can not fill in textures. The first approach is instead especially designed to copy texture parts. The exemplar-based inpainting algorithm needs no training of a model. It just uses the information of the image that should be inpainted. The FoE algorithm in contrast uses a trained model and the training is a time-consuming task. In my opinion the FoE may produce better results for a certain kind of images, since it is possible to train the model on images of this certain kind.

References

[1] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, "Image inpainting," in *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, ACM Press, 2000.

[2] A. Criminisi, P. Perez, and K. Toyama, "Region filling and object removal by exemplar-based image inpainting," *IEEE Transactions on Image Processing*, vol. 13, no. 9, pp. 1200–1212, 2004.

[3] S. Roth and M. Black, "Fields of experts," *International Journal of Computer Vision*, vol. 82, pp. 205–229, Apr. 2009.

[4] Z. Tauber, Z.-N. Li, and M. Drew, "Review and preview: Disocclusion by inpainting for image-based rendering.," *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, vol. 37, pp. 527–540, Jan. 2007.